Dynamic Decision Making Laboratory

Issue 15 Fall 2023

Carnegie Mellon University

Annual Summary from Coty — Founding Director, DDMLab

Inside this issue:

Dear Friends and Collaborators:

This is our annual update of the activities of the Dynamic Decision Making Laboratory (DDMLab) during the 2022-2023 academic year. Thanks to all of you who are interested in reading this update!

As it is a tradition, I will start my summary with news regarding research grants. In our last year's newsletter, we asked you, our readers, friends, and collaborators, to send your good vibes for the extensive engagement in work for a large proposal for an Institute for AI-enabled Societal Decision Making. Well, Thank you for the good vibes!!. It is my great pleasure to announce a new National Artificial Intelligence Research Institute: The Artificial Intelligence Institute for Societal Decision Making (AI-SDM) is funded by the **National Science Foundation (NSF)** and led by Carnegie Mellon University in collaboration with many other institutions, including: Harvard University, Howard University, Navajo Technical University, Penn State University, Texas A&M, University of Washington, and other organizations. I am a Research Co-Director together with Prof. Aarti Singh (Machine Learning Department at CMU). Under AI-SDM we expect to develop foundational research regarding human decision making under uncertainty, time constraints, and temporal dynamics and how to represent the cognitive decision-making process computationally. We also expect to demonstrate AI deployment in a number of challenging societal problems, such as disaster and emergency management and health resource allocation and interventions. A more complete story about this particularly noteworthy news, can be found in the CMU news archive.

We also continued to be supported with grants from many other institutions, including the **Air Force Office of Scientific Research (AFOSR)** for our work on the cognitive modeling of personalized and individualized recommender systems, in collaboration with researches at Aptima, Inc. **The Army Research Office (ARO)** Multi-University Research Initiative (MURI) program on cyber deception in collaboration with Harvard University. The **ARO US-Australia** International MURI (CATCH) on building cyber security teams of humans and bots (AI, machines). The Army **Research Laboratories (ARL)** program on collaborative research alliances on cybersecurity. Finally, the **Defense Advanced Research projects Agency (DARPA)** program on Artificial Social Intelligence for Successful Teams (ASIST), focused on collective intelligence and creating new models that can be used to collaborate with humans in teams to advance collective intelligence and adaptation.

As it will be clear through the research summaries of the DDMLab members in this newsletter, all these funding agencies are providing resources into advancing foundational research on dynamic decision making that has implications for a number of applied significant problems, notably: cybersecurity and human-AI interactions and collaborations.

In addition to the great news regarding funding for our work, the 2022-2023 academic year was particularly unique in many other significant ways:

Certificate of Completion

The Fulbright Foreign Scholarship Board and The Bureau of Educational and Cultural Affairs of the United States Department of State award this certificate to

Cleotilde Gonzalez

in recognition of successful completion of the

Fulbright Scholarship Program

May 2023
Washington, DC

Chair, Fulbright Foreign Scholarship Board

Assistant Secretary of State Bureau of Educational and Cultural Affairs

FULBRIGHT

I became a **Senior Fulbright Scholar** in Spain. I was affiliated faculty at the Department of Psychology at University of Malaga, during Spring 2023. This was my "sabbatical" semester and I engaged in a significant number of activities to bring together our research work at the DDMLab to various groups in Spain. In addition to my engagement with University of Malaga, I engaged in teaching activities with University Carlos III, in Madrid, and organized a significant exchange activity with Ellis Alicante.

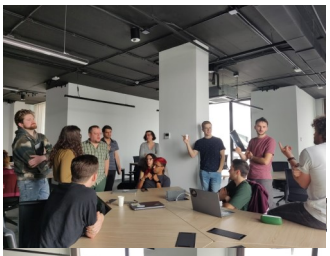

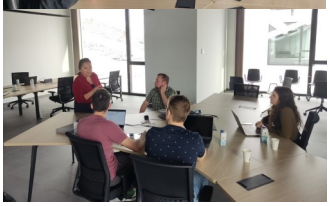

The ELLIS-Alicante and DDMLab workshop helped to generate new research ideas and collaborations. The workshop occurred during march 7-9, 2023 in Alicante, Spain. This was a collaboration between ELLIS, an institute of Human-Centered AI, located in Alicante, Spain, and the DDMLab.

Most incredible, is that I was able to invite a majority of my lab members to travel to Spain!! **Yes! I brought my lab to Spain to participate in this workshop.** Many Thanks to Nuria Oliver, Cristina Gonzalez and all the participants of the workshop for this amazing collaborative experience.

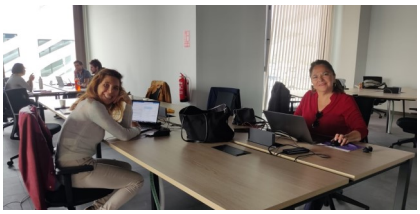

As it happens in cultural exchanges, what becomes most important is the *human* connections developed!! Our colleagues in Spain were incredible hosts and provided us with memories and experiences that we will never forget!! THANK YOU, Spain, Fulbright, University of Malaga, Ellis Alicante, and University Carlos III, Madrid. We will take you in our hearts forever.

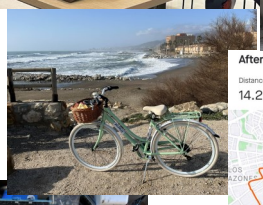

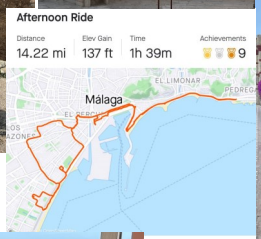

While in Spain, I got a new cruiser bicycle. I LOVED biking almost every day, next to the ocean, stopping for some tapas and a *caña*. My biking experience was unforgettable!

Back in Pittsburgh, I decided to get a similar cruiser bicycle. This one is the same color as my bike in Spain, but it is electric. I LOVE it!

In the past months, I had to take a little break from painting. But while in Spain I engaged in a new artistic activity: pottery!! My instructor Antonio was such a perfectionist, and had so much patience with me, that I was able to create some pieces that are actually pretty! (I think).

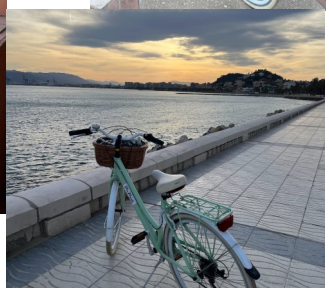

Overall, this academic year was VERY unique, enjoyable, memorable. I look forward to the 2023-2024 academic year and I hope you will continue to join us in making all our work count. ***Let's all work to make this a better world!!***

Coty

Farewell and Welcome New Members!

Farewells

We have to say good bye to two of our post-docs this year, Ngoc Nguyen and Baptiste Prebot, and one of our master student interns Tony Xi. Ngoc is in a tenure track Assistant Professor position in Computer Science at University of Dayton and Baptiste accepted a Human Factors specialist position at the Direction Generale de L'Armement, the research and development service of the French Department of Defense (Ministère des Armées) in Paris. Tony continued with his studies.

To everyone on this list, we are very proud and happy for all of you!! Good luck in your future endeavors!

Welcome!

Maria José Rodrigues Ferreira earned her Ph.D. in Computer Science and Engineering in 2023 from Instituto Superior Técnico, University of Lisbon (Portugal). Her research focuses on enhancing Social Agents' capabilities to assist people in a more personalized and engaging manner by exploring the intersection of Human-Agent Interaction and technology.

Roderick Seow received his Ph.D. in Cognitive Neuroscience in 2023 at Carnegie Mellon University. He is broadly interested in learning, generalization, and decision making, and has investigated these topics in the contexts of complex perceptual-motor skills and function learning. As a postdoctoral researcher with the Dynamic Decision Making Lab, he will explore topics surrounding human-AI collaboration.

Duncan Wood is a Ph.D. student in the Department of Social and Decision Sciences pursuing the Cognitive Decision Science track. The functional structure of memory, social information integration, social instance learning, and instance simulation are components of his research interests.

Research Updates from Lab Members During 2022-2023

From Tyler Malloy

During the past year I officially started my position as postdoctoral researcher in the DDMLab. Since then, I have begun working on projects related to human learning and decision making in abstract contexts, Human-AI collaboration in a variety of tasks, and optimal defense strategies against human attackers in a cybersecurity context.

The first project I began working on was related to a recently run experiment in CybOrg, which gave me a great opportunity to familiarize myself with the data collection methods of the DDMLab, as well as the cybersecurity environments we use. After this, I designed a simulated experiment to better understand the similarities and differences between Instance-Based Learning and a related approach that I have more experience with called Reinforcement Learning. This led to a project simulating the behavior of agents in a multiplayer game with an attack and defense role. While working on that project, I also began assisting Chase on his work in the minimap environment, which led him to design a model that utilizes a generative large language model to improve the training of an IBL agent in a simple gridworld and submitted a paper on that model to the AAAI symposium on the Integration of Cognitive Architectures and Generative Models.

In July, I presented at the IEEE Symposium on Security and Privacy workshop on Active Defense and Deception. This work was related to the simple defense task I simulated, which can be interesting to investigate due to a difficulty for AI to optimally defend against human attackers. To train the AI model, we first gave it experience in the role of the attacker. Our novel method was able to leverage this experience to improve its defense strategies, much in the same way that human defenders put themselves in the shoes of attackers. This was done by incorporating theory of mind reasoning with transfer of learning capabilities.

One major continuing project has been a collaboration with Yinuo and Coty with professor Fei Fang of the CMU Software and Societal System Department. This project focuses on how humans and machines transfer their learning from one context onto another, using a range of tasks from abstract to more complex. Recently our paper was accepted to the AAAI Conference on Human Computation and Crowdsourcing. We additionally submitted to the same AAAI symposium a paper on a model that compared how generative ML approaches can be integrated with IBL to improve transfer of learning. In the future, we hope to extend this work into investigating the complexities of how cognitive faculties like transfer of learning can impact decision making in tasks with multiple human or AI agents.

From Maria Ferreira

I earned my Ph.D. in Computer Science and Engineering from Instituto Superior Técnico, University of Lisbon (Portugal) in April of this year. My research centers around enhancing the capabilities of Social Agents to provide personalized and engaging assistance, achieved by exploring the intersection of Human-Agent Interaction and technology. My research interests encompass user-adapted environments, decision-making processes, and the manipulation of personality traits.

Prior to joining the DDMLab, I had the privilege of being invited to participate in their collaborative workshop with ELLIS Alicante in Spain in March. This three-day experience was truly enriching, allowing me to present my collaborative work on the effects of intergroup competition on the collective risk dilemma. Simultaneously, I teamed with members from both research groups, Tyler Malloy and Gergely D. Németh, on the proposal of a small project centered around a public transit system, employing a Stackelberg security game to model people's behaviors.

As a new member of the lab, I have delved into the practical applications of the lab's tools, such as PyIBL. My current focus has been on testing and refining established scenarios, including binary and 4-arm tasks. In this role, I've had the opportunity to employ various capabilities of PyIBL, such as partial matching and model tracing.

Building upon the insights gained from these tests, my upcoming objectives involve leveraging this knowledge to address a Phishing task. By doing so, we aim to create a more personalized approach to helping users adapt to dynamic scenarios.

Research Updates Continued

From Erin Bugbee

This past year marked my third at the DDMLab, where my doctoral work has focused on constructing computational cognitive models that illuminate how humans learn and make sequential decisions from their experiences. A core interest of mine is the optimal stopping problem, in which people must balance exploring options and collecting information with exploiting their knowledge and making a selection. I have developed IBL models that provide an explanation of how humans learn to make stopping decisions from experience. I have also shown how these models can accurately predict human stopping behavior in a variety of sequential decision tasks such as choosing when to purchase a flight ticket when prices change dynamically over time. To gain a deeper understanding of how various factors influence stopping behavior, I have developed a novel optimal stopping task. I will be conducting experiments with this task to gain insight into the impact of factors such as knowledge of the distribution of options and the type of feedback on stopping behavior.

Through a collaboration with Aptima as part of the TRUST'M project, I have investigated the strategies employed by intelligence analysts as they seek information to answer complex questions with the aid of a recommendation system. We have developed an IBL model that predicts the human decision regarding the recommended article, with the goal of cultivating trust between the human and the machine. This project bridges theoretical insights from cognitive science with human-machine teaming applications.

From giving a talk and exploring potential collaborations at the ELLIS-DDMLab Workshop in Alicante, Spain, to meeting leaders in the field at SJDM in San Diego, to presenting a poster at the 3rd Workshop on Mental Effort in Providence, RI, this year was truly enriching. Looking ahead, I have been named an Opportunity Scholar for posit::conf(2023) and am excited to connect with other data enthusiasts in Chicago. I have also had a poster accepted to SJDM in San Francisco and am awaiting the results for various other submissions.

Finally, this past summer I was an Applied Scientist Intern at Amazon Science on the Machine Learning University team, where I developed a workshop on fairness in large language models and taught various courses on topics in machine learning and generative AI.

This past year has been filled with opportunities for learning, connecting, and advancing my research, and I look forward to the year ahead in the DDMLab!

From Chase McDonald

Over the last year in the DDMLab, I've participated in a handful of projects that included human experimentation, cognitive modeling, and human-AI teaming. Under the ASIST grant, I continued work on a project originally led by former DDMLab member Ngoc Nguyen. In it, we use the Team Minimap Search and Rescue task with four human participants and manipulate the presentation of measures that capture the collective intelligence of the team in the task. Our work, recently submitted to the ACM Collective Intelligence conference, investigates how these presentations impact team performance and behavior.

In recent work, accepted to the 2023 AAAI Fall Symposium on the Integration of Cognitive Architectures and Generative Models, we explored how advances in large language models (LLMs) can be utilized in cognitive modeling. Specifically, we explored the use of LLMs to interpret task instructions and produce both steps required to complete a task and the value of reaching each step, then used these to construct intrinsic rewards. Using the LLM-derived intrinsic rewards, we were able to significantly improve IBL model learning. Future work will both compare the output of the LLM to human participants, as well as scale up the task and model complexity.

A major focus over the last year has been the development of an experimental framework and a task to use for human-AI teaming. To the former, we have developed a general framework that takes simulation environments (e.g., gymnasium environments) and allows them to be used in online interactive experiments with humans and AI. To use this framework, we've developed a new experimental task that allows us to construct gridworld-based experiments, train cognitive models and reinforcement learning agents, and test human-AI teaming in cooperative settings. As we look forward to the next year, we'll focus on developing models and designing experiments that incorporate insights from cognitive modeling to improve the collaborative capacity of AI.

Research Updates Continued

From Yinuo Du

During the past year, I kept working on cyber deception by building a high-fidelity platform to empirically evaluate the effectiveness of cyber deception techniques and game-theoretic deceptive defense strategies. In addition, I worked on human-AI teaming in the context of cyber security with Baptiste and Tyler. In it, we use the Team Defense Game of a human-agent dyad and manipulate the underpinning algorithm of the autonomy agent. This work is in preparation for submission and will be presented at INFORMS'23. I was also lucky to be involved in the projects and submissions led by Tyler on transfer of learning.

I was fortunate to get the opportunities to present a poster at HFES'22, give talks at Cylab partners conference'22, HICSS'23 conference in Hawaii, ELLIS-DDM Lab Workshop, and MURI workshop. Under grant ARL-CRA, I gave live demos of game -theoretic deceptive defense at the bootcamps. Thanks to student scholarship from WiCyS, I went to Denver and met with inspiring women in cyber security.

A major focus of the last year has been cyber defense with human-like cognitive agents, RL-based autonomy agents, and collaborative defense of human-AI teams. In the coming academic year, I will further my exploration in these areas and try to build agents that can facilitate, teach, and cooperate with humans and can be applied in more general domains.

From Roderick Seow

In my dissertation, I developed a novel theory of rule discovery in function learning in order to account for how people could discover and apply linear rules from past observations of pairs of variable values. As an incoming member of the Dynamic Decision Making Lab, I look forward to exploring the cognitive processes that underlie decision making in human-AI collaborative contexts. I am particularly interested in how these processes compare to those of single-agent and human-human contexts, and how biases, decisions, and feedback interact to affect the learning and performance trajectories of such collaborations.

From Duncan Wood

As a new SDS PhD student, I'm excited to start working on human-machine collaboration with the DDMLab. Coming from Rutgers University, where I studied cognitive science and economics, my previous research used computational models for a wide range of tasks, including automated detection of rail trespassing and agent-based models of homophilic friendship formation. My economics honors thesis used agent-based models to reconcile disagreements in theoretical and empirical findings by populating a market with sellers using reinforcement learning and buyers using cognitively informed decision rules. I'm looking forward to delving deeper into cognition and dynamic decision-making environments with my future modeling research with the DDMLab.

From Don Morrison

A new version of PyIBL, 5.0, was released. The primary changes are that it is substantially faster for some large models, details of the API have been changed to be more flexible for the future, and it now requires at least Python 3.8. Details of the changes are in the PyIBL documentation at http://pyibl.ddmlab.com/.

From Jeffrey Flagg

This spring, I had the great pleasure to help plan our workshop with ELLIS-Alicante. I would like to thank Scientific Director and Co-founder of ELLIS Alicante, Nuria Oliver; Office manager, Cristina Gonzalez; and PhD student, Aditya Gulati for all their hard work. During this 3-day workshop, researchers worked to encourage and develop novel collaborations in the study of human-AI collaborations and the role of AI in human societies.

I am also helping to plan a workshop this fall with members of the Army Research Office. This workshop, titled The Future of Cyber Deception, will bring together leading researchers from cyber security, AI, as well as cognitive scientists, who work on cyber deception against adaptive and intelligent adversaries. This workshop will shed light on key challenges and interdisciplinary research opportunities, with the ultimate goal of improving cyber defence.

I have also continued assisting on various experiments, helping test materials, reviewing papers, maintaining our website and social media materials, and integrating new members. I have had the pleasure to help review several paper submissions for DDMLab members. I also have continued to serve on CMU's OSF Advisory Board which helps promote best practices in the greater CMU community.

Research Updates Continued

From Ngoc Nguyen

This year, I was involved in human-machine teaming projects with Coty and other lab members. Together with Chase, Coty, and Anita Woolley's group, we have studied real-time collaborative process metrics to measure and predict collective intelligence in teamwork. The findings of this study are anticipated to guide interventions to improve team performance. In the joint work with Chase, Tyler, and Coty, we have explored the use of large language models as a reward function for improving computational cognitive modeling. This research has recently been accepted for publication at the AAAI 2023 Fall Symposium on Integration of Cognitive Architectures and Generative Models.

Regarding publications, we had success in publishing our work on Minimap, a dynamic interactive environment for studying dynamic decision making in search and rescue missions, in the Behavior Research Method journal. Together with Nhat and Coty, we have published our work on Multi-agent Instance-based Learning (MAIBL) algorithms. These algorithms are a combination of cognitive mechanisms of IBLT and the techniques of multi-agent deep reinforcement learning models, which are designed to tackle coordinated multi-agent systems in stochastic environments from the perspective of independent learners. This work has been published in the ACM Transactions on Autonomous and Adaptive Systems.

After three wonderful years at DDM Lab, I have accepted a position as an Assistant Professor in Computer Science at the University of Dayton. I am deeply grateful to Coty for her guidance and support. I feel fortunate to have been a part of the DDM Lab and appreciate the opportunities and experiences it has provided me during my postdoc. I look forward to continuing collaboration with DDM Lab members in the future.

From Baptiste Prebot

During my short but wonderful time in the DDMLab, I worked mainly on human-artificial intelligence collaboration in cyber defense. We built a framework and developed two cybersecurity games to study humans, cognitive models and their interactions as teammates and adversaries in cyber defense scenarios.

Although I'm now back in France, I still enjoy being involved in the project, helping to design future experiments that will hopefully test a truly collaborative and adaptive cognitive model. I am now a human factors specialist at the French Ministry of Defense (Ministère des Armées), where I oversee human factors and ergonomic aspects in the development of command and control systems for the French Space Command.

Fireworks at CMU: 4th of July, 2023

A trip to Fallingwater: August, 2023

ELLIS-Alicante / DDMLab 2023 Workshop

ELLIS-Alicante in Spain and the DDMLab joined forces this year to explore research collaborations on Human-Centered Artificial Intelligence. In the 3-day workshop researchers from both laboratories came together to encourage and develop novel collaborations in the study of human-AI complementarity and the role of AI in human societies. The goal of this workshop was to generate new research ideas that will advance the foundational research of human-AI collaborations and interdependencies.

Many thanks to all the people that made the 2023 ELLIS-Alicante / DDMLab workshop possible, especially ELLIS co-founder Nuria Oliver, Office Manager Cristina Gonzalez, and PhD student Aditya Gulati.

Recent Publications

In the past year, we published several journal articles and chapters authored by members of the DDMLab and our collaborators. For a full list of publications, please see the publications page on the laboratory's website at https://www.cmu.edu/dietrich/sds/ddmlab/publications.html

Some Recent Publications Highlights:

Gonzalez, C., Admoni, H., Brown, S., & Woolley, A.W. (2023), COHUMAIN: Building the Socio-Cognitive Architecture of Collective Human–Machine Intelligence. Topics in Cognitive Science. https://doi.org/10.1111/tops.12673.

Gonzalez, C., Aggarwal, P., Cranford, E. A., Lebiere, C. (2023). Adaptive Cyberdefense with Deception: A Human-AI Cognitive Approach. In T. Bao et al., (eds.). Cyber Deception, Techniques, Strategies, and Human Aspects. Vol. 89. pp. 41-57, Springer. https://doi.org/10.1007/978-3-031-16613-6_3.

Du, Y., Prebot, B., Xi, X. & Gonzalez, C., (2023). A Cyber-War Between Bots: Human-Like Attackers are More Challenging for Defenders than Deterministic Attackers. In Proceedings of the 56 Hawaii International Conference on System Sciences. (pp. 856-865). HICSS 2023, January 7-10, 2023, Maui, HI. https://hdl.handle.net/10125/102736

Galesic, M., Barkoczi, D., Berdahl, A., Biro, D., Carbone, G., Gonzalez, C., Kandler, A., Kao, A., Kendal, R., Kline, M., Lee, E., Massari, G. F., Mesoudi, A., Olsson, H., Pescetelli, N., Sloman, S. J., Smaldimo, P. E., & Stein, D. L. (2023). Beyond collective intelligence: Collective adaptation. J. R. Soc. Interface. 20: 20220736. https://doi.org/10.1098/rsif.2022.0736

Malloy, T., & Gonzalez, C. (2023). Learning to Defend by Attacking (and Vice-Versa): Transfer of Learning in Cybersecurity Games. 2nd Workshop on Active Defense and deception (EuropS&P'23 AD&D Workshop). Delft, Netherlands. July 3-7, 2023. https://adnd23.hotcrp.com/doc/adnd23-final1.pdf?cap=hcav1CWPgLKtDzWtVbRaRBHhRpRpB.

Nguyen, T. N., & Gonzalez, C. (2023). Minimap: An Interactive Dynamic Decision Making Game for Search and Rescue Missions. Behavioral Research Methods. In press.

Singh, K., Aggarwal, P., Rajivan, P., & Gonzalez, C. (2023). Cognitive elements of learning and discriminability in anti-phishing training. Computers & Security. 103105. https://doi.org/10.1016/j.cose.2023.103105.

Zhao, M., Eadeh, F. R., Nguyen, T. N., Gupta, P., Admoni, H., Gonzalez, C., & Woolley, A. (in press). Teaching Agents to Understand Teamwork: Evaluating and Predicting Collective Intelligence as a Latent Variable via Hidden Markov Models. Computers in Human Behavior. Volume 139, February 2023, 107524. https://doi.org/10.1016/j.chb.2022.107524.

Highlighted Projects and Events

This year, the DDMLab was a part of a major new grant! The National Science Foundation grant will help fund the AI Institute for Societal Decision Making. The AI-SDM will improve the responses to societal challenges such as disaster management and public health by creating human-centric AI tools to assist with critical decisions. The institute will also develop interdisciplinary training to bolster effective and rapid response in uncertain and dynamic situations. Coty will serve as the institute's research co-director!

AI-SDM is the fifth NSF-funded AI institute to include researchers from CMU, and the first to be led by the university's expertise. CMU faculty already contribute to the AI Institute for Collaborative Assistance and Responsive Interaction for Networked Groups (AI-CARING), the AI Institute for Future Edge Networks and Distributed Intelligence (AI-EDGE), the USDA-NIFA AI Institute for Resilient Agriculture (AIIRA) and the Institute for Agricultural AI for Transforming Workforce and Decision Support (AgAID). These institutes were established in 2021. For information about NSF institutes, visit:

https://new.nsf.gov/funding/opportunities/national-artificial-intelligence-research

Please view the AI Institute for Societal Decision Making's website here! **https://www.cmu.edu/ai-sdm/**

Adventures During 2022 - 2023

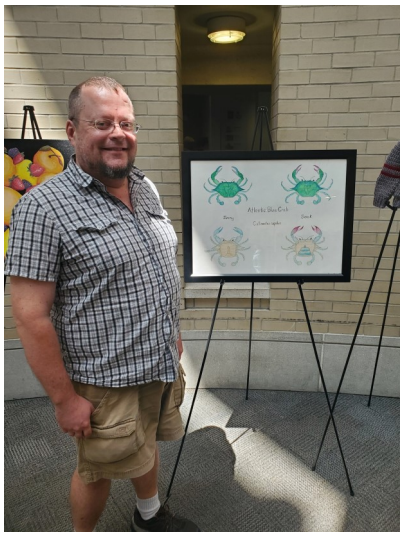

Top: Roderick Seow, Coty Gonzalez, Jeffrey Flagg, Duncan Wood, Tyler Malloy, Yinuo Du, Maria Ferreira, and Chase McDonald at Fallingwater.

Left: DDMLab member Jeffrey Flagg had art displayed at this year's CMU Community Picnic.

Right: Baptiste Prebot, Chase McDonald, and Erin Bugbee at the Santa Barara Castle in Alicante, Spain.

Down: Coty with two students she met in Sydney, during CogSci 2023

Dynamic Decision Making Laboratory

DDMLab-Porter Hall
Carnegie Mellon University
5000 Forbes Avenue
Social and Decision Making Sciences
Porter Hall 223-G
Pittsburgh, PA 15213
e-mail: ddmlab(at)andrew.cmu.edu
Phone: 412-268-9547
412-268-8439

Thank you for your interest in our

2023 Newsletter!

Keep up with us on our webpage or follow us on

Facebook LinkedIn twitter or youtube